

Infinite matroids in graphs

Henning Bruhn

Reinhard Diestel

Abstract

It has recently been shown that infinite matroids can be axiomatized in a way that is very similar to finite matroids and permits duality. This was previously thought impossible, since finitary infinite matroids must have non-finitary duals.

In this paper we illustrate the new theory by exhibiting its implications for the cycle and bond matroids of infinite graphs. We also describe their algebraic cycle matroids, those whose circuits are the finite cycles and double rays, and determine their duals. Finally, we give a sufficient condition for a matroid to be representable in a sense adapted to infinite matroids. Which graphic matroids are representable in this sense remains an open question.

1 Introduction

In the current literature on matroids, infinite matroids are usually ignored, or else defined like finite ones¹ with the following additional axiom:

- (I4) An infinite set is independent as soon as all its finite subsets are independent.

We shall call such set systems *finitary matroids*.

The additional axiom (I4) reflects the notion of linear independence in vector spaces, and also the absence of (finite) circuits from a set of edges in a graph. More generally, it is a direct consequence of (I4) that circuits, defined as minimal dependent sets, are finite.

An important and regrettable consequence of the additional axiom (I4) is that it spoils duality, one of the key features of finite matroid theory. For example, the cocircuits of an infinite uniform matroid of rank k would be the sets missing exactly $k - 1$ points; since these sets are infinite, however, they cannot be the circuits of another finitary matroid. Similarly, every bond of an infinite graph would be a circuit in any dual of its cycle matroid—a set of edges minimal with the property of containing an edge from every spanning tree—but these sets can be infinite and hence will not be the circuits of a finitary matroid.

¹The augmentation axiom is required only for finite sets: given independent sets I, I' with $|I| < |I'| < \infty$, there is an $x \in I' \setminus I$ such that $I + x$ is again independent.

This situation prompted Rado in 1966 to ask for the development of a theory of non-finitary infinite matroids with duality [15, Problem P531]. In the late 1960s and 70s, a number of such theories were proposed; see [3] for references. One of these, the ‘B-matroids’ proposed by Higgs [12], were later shown by Oxley [13] to describe the models of any theory of infinite matroids that admitted both duality and minors as we know them. However, Higgs did not present his ‘B-matroids’ in terms of axioms similar to those for finite matroids. As a consequence, theorems about finite matroids whose proofs rested on these axioms could not be readily extended to infinite matroids, even when this might have been possible in principle.

With the axioms from [3] presented in the next section, this could now change: it should be possible now to extend many more results about finite matroids to infinite matroids, either by

- adapting their proofs based on the finite axioms to the (very similar) new infinite axioms,

or by

- finding a sequence of finite matroids that has the given infinite matroid as a limit, and is chosen in such a way that the instances of the theorem known for those finite matroids imply a corresponding assertion for the limit matroid.

After presenting our new axioms in Section 2, we apply them in Sections 3 and 4 to see what they mean for graphs. We shall see that, for matroids whose circuits are the (usual finite) cycles of a graph, our axioms preserve what would be wrecked by the finitary axiom (I4): that their duals are the matroids whose circuits are the bonds of our graph – even though these can now be infinite.

The converse is also nice. The dual M^* of the (finitary) matroid M whose circuits are the finite bonds of an infinite graph cannot be finitary; indeed, trivial exceptions aside, the duals of finitary matroids are never finitary [19, 1]; see also [5]. So M^* will have to have infinite circuits.² Excitingly, these circuits turn out to be familiar objects: when the graph is locally finite, they are the edge sets of the topological circles in its Freudenthal compactification, which already have a fixed place in infinite graph theory quite independently of matroids [8].

We shall see further that, for planar graphs, matroid duality is now fully compatible with graph duality as explored in [2]. And Whitney’s theorem, that a graph is planar if and only if its cycle matroid has a graphic dual, now has an infinite version too.

In some infinite graphs G , including all locally finite ones, the elementary algebraic cycles³ form the circuits of a matroid, the *algebraic cycle matroid* of G . We introduce this matroid, which was already studied by Higgs [11], in Section 3. In Section 5, we determine its dual.

²We are using here that every dependent set contains a minimal such. This is indeed true.

³These are the minimal 1-chains (possibly infinite) with zero boundary: the edge sets of finite cycles and of 2-way infinite paths.

In Section 6, finally, we consider representability. Infinite matroids that are representable in the usual sense are finitary, and hence cannot have representable duals. We suggest an adapted notion of representability based on infinite sums of functions to a field, and establish a sufficient condition for when this defines a matroid. Examples include the algebraic cycle matroids of graphs introduced in Section 3, and the algebraic cycle matroids of higher-dimensional complexes [3].

2 Axioms

We now present our five sets of axioms for finite or infinite matroids, in terms of independent sets, bases, circuits, closure and rank. These axioms were first stated in [3], and proved to be equivalent to each other in the usual sense. For finite or finitary matroids they coincide with the usual finite matroid axioms.

Let E be any set, finite or infinite; it will be the default ground set for all matroids considered in this paper. We write 2^E for its power set. The set of all pairs (A, B) such that $B \subseteq A \subseteq E$ will be denoted by $(2^E \times 2^E)_{\subseteq}$; for its elements we usually write $(A|B)$ instead of (A, B) . Unless otherwise mentioned, the terms ‘minimal’ and ‘maximal’ refer to set inclusion. Given $\mathcal{E} \subseteq 2^E$, we write \mathcal{E}^{\max} for the set of maximal elements of \mathcal{E} , and $[\mathcal{E}]$ for the *down-closure* of \mathcal{E} , the set of subsets of elements of \mathcal{E} . For $F \subseteq E$ and $x \in E$, we abbreviate $F \setminus \{x\}$ to $F - x$ and $F \cup \{x\}$ to $F + x$. We shall not distinguish between infinite cardinalities and denote all these by ∞ ; in particular, we shall write $|A| = |B|$ for any two infinite sets A and B . The set \mathbb{N} contains 0.

One central axiom that features in all our axiom systems is that every independent set extends to a maximal one, even inside any restriction $X \subseteq E$.⁴ The notion of what constitutes an independent set, however, will depend on the type of axioms under consideration. We therefore state this extension axiom in a more general form first, without reference to independence, so as to be able to refer to it later from within different contexts.

Let $\mathcal{I} \subseteq 2^E$. The following statement describes a possible property of \mathcal{I} .

- (M) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}$, the set $\{I' \in \mathcal{I} \mid I \subseteq I' \subseteq X\}$ has a maximal element.

Note that the maximal superset of I in $\mathcal{I} \cap 2^X$ whose existence is asserted in (M) need not lie in \mathcal{I}^{\max} .

2.1 Independence axioms

The following statements about a subset \mathcal{I} of 2^E are our *independence axioms*:

- (I1) $\emptyset \in \mathcal{I}$.
(I2) $[\mathcal{I}] = \mathcal{I}$, that is, \mathcal{I} is closed under taking subsets.

⁴Interestingly, we shall not need to require that every dependent set contains a minimal dependent set. We need that too, but it follows [3].

(I3) For all $I \in \mathcal{I} \setminus \mathcal{I}^{\max}$ and $I' \in \mathcal{I}^{\max}$, there is an $x \in I' \setminus I$ such that $I+x \in \mathcal{I}$.

(IM) \mathcal{I} satisfies (M).

When a set $\mathcal{I} \subseteq 2^E$ satisfies the independence axioms, we call the pair (E, \mathcal{I}) a *matroid* on E . We call every element of \mathcal{I} an *independent* set, every element of $2^E \setminus \mathcal{I}$ a *dependent* set, the maximal independent sets *bases*, and the minimal dependent sets *circuits*. This matroid is *finitary* if it also satisfies (I4) from the Introduction, which is equivalent to requiring that every circuit be finite [3].

The $2^E \rightarrow 2^E$ function mapping a set $X \subseteq E$ to the set

$$\text{cl}(X) := X \cup \{x \mid \exists I \subseteq X : I \in \mathcal{I} \text{ but } I+x \notin \mathcal{I}\}$$

will be called the *closure operator* on 2^E associated with \mathcal{I} .

The $(2^E \times 2^E)_{\subseteq} \rightarrow \mathbb{N} \cup \{\infty\}$ function r that maps a pair $A \supseteq B$ of subsets of E to

$$r(A|B) := \max \{|I \setminus J| : I \supseteq J, I \in \mathcal{I} \cap 2^A, J \text{ maximal in } \mathcal{I} \cap 2^B\}$$

will be called the *relative rank function* on the subsets of E associated with \mathcal{I} . This maximum is always attained, and independent of the choice of J [3].

2.2 Basis axioms

The following statements about a set $\mathcal{B} \subseteq 2^E$ are our *basis axioms*:

(B1) $\mathcal{B} \neq \emptyset$.

(B2) Whenever $B_1, B_2 \in \mathcal{B}$ and $x \in B_1 \setminus B_2$, there is an element y of $B_2 \setminus B_1$ such that $(B_1 - x) + y \in \mathcal{B}$.

(BM) The set $\mathcal{I} := [\mathcal{B}]$ of all \mathcal{B} -independent sets satisfies (M).

2.3 Closure axioms

The following statements about a function $\text{cl}: 2^E \rightarrow 2^E$ are our *closure axioms*:

(CL1) For all $X \subseteq E$ we have $X \subseteq \text{cl}(X)$.

(CL2) For all $X \subseteq Y \subseteq E$ we have $\text{cl}(X) \subseteq \text{cl}(Y)$.

(CL3) For all $X \subseteq E$ we have $\text{cl}(\text{cl}(X)) = \text{cl}(X)$.

(CL4) For all $Z \subseteq E$ and $x, y \in E$, if $y \in \text{cl}(Z+x) \setminus \text{cl}(Z)$ then $x \in \text{cl}(Z+y)$.

(CLM) The set \mathcal{I} of all *cl-independent* sets satisfies (M). These are the sets $I \subseteq E$ such that $x \notin \text{cl}(I-x)$ for all $x \in I$.

2.4 Circuit axioms

The following statements about a set $\mathcal{C} \subseteq 2^E$ are our *circuit axioms*:

- (C1) $\emptyset \notin \mathcal{C}$.
- (C2) No element of \mathcal{C} is a subset of another.
- (C3) Whenever $X \subseteq C \in \mathcal{C}$ and $(C_x \mid x \in X)$ is a family of elements of \mathcal{C} such that $x \in C_y \Leftrightarrow x = y$ for all $x, y \in X$, then for every $z \in C \setminus (\bigcup_{x \in X} C_x)$ there exists an element $C' \in \mathcal{C}$ such that $z \in C' \subseteq (C \cup \bigcup_{x \in X} C_x) \setminus X$.
- (CM) The set \mathcal{I} of all \mathcal{C} -independent sets satisfies (M). These are the sets $I \subseteq E$ such that $C \not\subseteq I$ for all $C \in \mathcal{C}$.

Axiom (C3) coincides for $|X| = 1$ with the usual ('strong') circuit elimination axiom for finite matroids. In particular, it implies that adding an element to a basis creates at most one circuit; the fact that it does create such a (*fundamental*) circuit is trivial when bases are defined from these circuit axioms (as maximal sets not containing a circuit), while if we start from the independence axioms it follows from the fact, mentioned before, that every dependent set contains a minimal one [3]. We remark that the usual finite circuit elimination axiom is too weak to guarantee a matroid [3].

2.5 Rank axioms

The following statements about a function $r: (2^E \times 2^E)_{\subseteq} \rightarrow \mathbb{N} \cup \{\infty\}$ are our (relative) *rank axioms*:

- (R1) For all $B \subseteq A \subseteq E$ we have $r(A|B) \leq |A \setminus B|$.
- (R2) For all $A, B \subseteq E$ we have $r(A|A \cap B) \geq r(A \cup B|B)$.
- (R3) For all $C \subseteq B \subseteq A \subseteq E$ we have $r(A|C) = r(A|B) + r(B|C)$.
- (R4) For all families (A_γ) and B such that $B \subseteq A_\gamma \subseteq E$ and $r(A_\gamma|B) = 0$ for all γ , we have $r(A|B) = 0$ for $A := \bigcup_\gamma A_\gamma$.
- (RM) The set \mathcal{I} of all r -independent sets satisfies (M). These are the sets $I \subseteq E$ such that $r(I|I - x) > 0$ for all $x \in I$.

For finite matroids, these axioms (with (R4) and (RM) becoming redundant) are easily seen to be tantamount to the usual axioms for an absolute rank function R derived as $R(A) := r(A|\emptyset)$, or conversely with $r(A|B) := R(A) - R(B)$ for $B \subseteq A$.

3 Bond and cycle matroids

In this section we develop the theory of our axioms to see what it yields for the usual matroids for graphs when these are infinite. See [3] for applications to other structures than graphs. All our graphs may have parallel edges and loops.

A well-known matroid associated with a finite graph G is its cycle matroid: the matroid whose circuits are the edge sets of the cycles in G . The bases of this matroid are the edge sets of the *spanning forests* of G , the sets that form a spanning tree in every component of G . This construction works in infinite graphs too: the edge sets of the finite cycles in G form the circuits of a finitary matroid $M_{\text{FC}}(G)$, whose bases are the edge sets of the spanning forests of G . We shall call $M_{\text{FC}}(G)$ the *finite-cycle matroid* of G . Similarly, we let the *finite-bond matroid* $M_{\text{FB}}(G)$ of G be the matroid whose circuits are the finite bonds of G . (A *bond* is a minimal non-empty cut.) This, too, is a finitary matroid.

If G is finite, then $M_{\text{FC}}(G)$ and $M_{\text{FB}}(G)$ are dual to each other. For infinite G , however, things are different. As remarked earlier, the duals of finitary matroids are not normally finitary, so the duals of $M_{\text{FC}}(G)$ and $M_{\text{FB}}(G)$ will in general have infinite circuits. In the case of $M_{\text{FC}}(G)$, its cocircuits are the expected ones, the (finite or infinite) bonds:

Theorem 1. *Let G be any graph.*

- (i) *The bonds of G , finite or infinite, are the circuits of a matroid $M_{\text{B}}(G)$.*
- (ii) *This matroid is the dual of the finite-cycle matroid $M_{\text{FC}}(G)$ of G .*

The matroid $M_{\text{B}}(G)$ defined in Theorem 1 will be called the *bond matroid* of G . We defer the proof of the theorem to Section 4; it is essentially the same as for finite graphs, although now the bonds can be infinite.

Similarly, the dual of $M_{\text{FB}}(G)$ will in general have infinite circuits. Ideally, these would form some sort of ‘infinite cycles’ in G . ‘Infinite cycles’ have indeed been considered before for graphs, though in a purely graph-theoretic context: there is a topological such notion that makes it possible to extend classical results about cycles in finite graphs (such as Hamilton cycles) to infinite graphs, see [8] and [17] in this issue. Rather strikingly, it turns out that these ‘infinite cycles’ are the solution also to our problem: their edge sets are precisely the (possibly infinite) cocircuits of $M_{\text{FB}}(G)$.

In order to define those ‘infinite cycles’, we need to endow our given graph G with a topology. A *ray* is a one-way infinite path. Two rays are *edge-equivalent* if for any finite set F of edges there is a component of $G - F$ that contains subrays of both rays. The equivalence classes of this relation are the *edge-ends* of G ; we denote the set of these edge-ends by $\mathcal{E}(G)$.

Let us view the edges of G as disjoint topological copies of $[0, 1]$, and let X_G be the quotient space obtained by identifying these copies at their common vertices. The set of inner points of an edge e will be denoted by \dot{e} . We now define a topological space $\|G\|$ on the point set of $X_G \cup \mathcal{E}(G)$ by taking as our open sets the unions of sets \bar{C} , where C is a connected component of $X_G - Z$

for some finite set $Z \subset X_G$ of inner points of edges, and \tilde{C} is obtained from C by adding all the edge-ends represented by a ray in C .⁵

When G is connected, $\|G\|$ is a compact topological space [16], although in general it need not be Hausdorff: the common starting vertex of infinitely many otherwise disjoint equivalent rays, for example, cannot be distinguished topologically from the edge-end which those rays represent, and neither can two vertices joined by infinitely many edges or edge-disjoint paths.⁶ However if G is locally finite, then $\|G\|$ coincides with the (Hausdorff) Freudenthal compactification of G . See Section 4 for more properties of $\|G\|$.

For any set $X \subseteq \|G\|$ we call

$$E(X) := \{e \in E(G) : \dot{e} \subseteq X\}$$

the *edge set* of X . A subspace C of $\|G\|$ that is homeomorphic to S^1 is a *circle* in $\|G\|$. One can show that $\bigcup E(C)$ is dense in C , so C lies in the closure of the subgraph formed by its edges [16]. In particular, there are no circles consisting only of edge-ends.

A subspace $X \subseteq \|G\|$ is a *standard subspace* if it is the closure in $\|G\|$ of a subgraph of G . A *topological spanning tree* of G is a standard subspace T of $\|G\|$ that is path-connected and contains $V(G)$ but contains no circle. Note that, since standard subspaces are closed, T will also contain $\mathcal{E}(G)$.

Theorem 2. *Let G be any connected⁷ graph.*

- (i) *The edge sets of the circles in $\|G\|$ are the circuits of a matroid $M_C(G)$, the cycle matroid of G .*
- (ii) *The bases of $M_C(G)$ are the edge sets of the topological spanning trees of G .*
- (iii) *The cycle matroid $M_C(G)$ is the dual of the finite-bond matroid $M_{FB}(G)$.*

We shall prove Theorem 2 in Section 4.

In the finite world, matroid duality is compatible with graph duality in that the dual of the cycle matroid of a finite planar graph G is the cycle matroid of its (geometric or algebraic) dual G^* . Duality for infinite graphs has come to be properly understood only recently [2]. But now that we have matroid duality as well, it turns out that the two are again compatible. In the remainder of this section we briefly explain how infinite graph duality is defined, and then show its compatibility with matroid duality.

When one tries to define abstract graph duality so that it satisfies the minimum requirement of capturing the geometric duality of locally finite graphs in the plane (where one has a dual vertex for every face and a dual edge between

⁵Note that $\|G\|$ induces on G a topology coarser than that of X_G . We needed X_G only provisionally, so that the connected components used in the definition of $\|G\|$ were defined.

⁶A consequence that takes some getting used to is that circles in $\|G\|$ need not look ‘round’: an edge joining two indistinguishable vertices, for example, will form a topological loop with either one of them.

⁷The theorem extends to disconnected graphs in the obvious way.

vertices representing two faces for every edge that lies on the boundary of both these faces), the first thing one realizes is that by taking duals one will leave the class of locally finite graphs: the dual of a ray, for example, is a vertex with infinitely many loops. On the other hand, Thomassen [18] showed that any class of graphs for which duality can be reasonably defined cannot be much larger: these graphs have to be *finitely separable* in that every two vertices can be separated by finitely many edges.⁸

It was finally shown in [2] that the class of finitely separable graphs is indeed the right setting for infinite graph duality, defined as follows. Let G be a finitely separable graph. A graph G^* is called a *dual of G* if there is a bijection

$$*: E(G) \rightarrow E(G^*)$$

such that a set $F \subseteq E(G)$ is the edge set of a circle in $\|G\|$ if and only if $F^* := \{e^* \mid e \in F\}$ is a bond of G^* .⁹ Duals defined in this way behave just as for finite graphs:

Theorem 3. [2] *Let G be a countable finitely separable graph.*

- (i) *G has a dual if and only if G is planar.*
- (ii) *If G^* is a dual of G , then G^* is finitely separable, G is a dual of G^* , and this is witnessed by the inverse bijection of $*$.*
- (iii) *Duals of 3-connected graphs are unique, up to isomorphism.*

At the time, the reason for defining graph duality as above was purely graph-theoretic: it appeared (and still appears) to be the unique way to make all three statements of Theorem 3 true for infinite graphs. As matroid duality was developed independently of graph duality, it is thus remarkable—and adds to the justification of both notions—that the two are once more compatible, as far as remains possible in an infinite setup:

Theorem 4. *Let G and G^* be a pair of countable dual graphs, each finitely separable, and defined on the same edge set E . Then*

$$M_{\text{FB}}(G) = M_{\text{C}}^*(G) = M_{\text{B}}^*(G^*) = M_{\text{FC}}(G^*).$$

Proof. The first equality is Theorem 2 (iii). The last equality is Theorem 1 (ii) (after dualizing). The middle equality follows from $M_{\text{C}}(G) = M_{\text{B}}(G^*)$, which is a direct consequence of the definition of a dual graph. \square

Finally, we obtain an infinite analogue of Whitney’s theorem that a finite graph is planar if and only if the dual of its cycle matroid is ‘graphic’, i.e., is the cycle matroid of another finite graph. In our more general context, let us call a matroid *graphic* if it is isomorphic to the cycle matroid of some graph, and *finitely graphic* if it is isomorphic to the finite-cycle matroid of some graph.

⁸Christian, Richter and Rooney [7] define certain dual objects for arbitrary planar graphs; however these objects are ‘graph-like spaces’, not graphs.

⁹We are cheating a bit here, but only slightly. In [2], these circles are taken not in $\|G\|$ but in a slightly different space \tilde{G} . However, while the circles in \tilde{G} may differ slightly from those in $\|G\|$, their edge sets are the same. This is not hard to see directly; it also follows from Theorems 6.3 and 6.5 in [10] in conjunction with Satz 4.3 and 4.5 in [16].

Theorem 5. *The following three assertions are equivalent for a countable finitely separable graph G :*

- (i) G is planar;
- (ii) $M_C^*(G)$ is finitely graphic;
- (iii) $M_{FC}^*(G)$ is graphic.

We shall prove Theorem 5 in Section 4. The equivalence of (i) and (ii) can also be derived from a more general result of Christian et al. on ‘graph-like spaces’ [7].

Another natural matroid in a locally finite graph G is its algebraic cycle matroid: the matroid whose circuits are the *elementary algebraic cycles* of G , the minimal non-empty edge sets inducing even degrees at all the vertices. Clearly, these are the edge sets of the finite cycles in G and those of its *double rays*, its 2-way infinite paths.

The elementary algebraic cycles do not form a matroid in every infinite graph: it is easy to check [3, Section 6] that they do not satisfy our circuit axioms when G is the *Bean graph* shown in Figure 1. However, by a result of Higgs [11] made applicable to our matroids by [3, Theorem 5.1], this is essentially the only counterexample:




Figure 1: The Bean graph

Theorem 6 (Higgs 1969). *The elementary algebraic cycles of an infinite graph G are the circuits of a matroid on its edge set $E(G)$ if and only if G contains no subdivision of the Bean graph.*

Corollary 7. *The elementary algebraic cycles of any locally finite graph are the circuits of a matroid.* \square

We call the matroid from Theorem 6 the *algebraic cycle matroid* $M_{AC}(G)$ of the graph G . In Section 5 we determine its dual $M_{AC}^*(G)$: it is the matroid whose circuits are the minimal non-empty cuts of G at least one side of which contains no ray.

The algebraic cycle matroid is *representable* in a sense adapted to non-finitary matroids, which we discuss in Section 6. We shall also prove a general sufficient condition for this notion of representability.

4 Proofs of Theorems 1, 2 and 5

We begin with the easy proof of Theorem 1, which we restate:

Theorem 1. *Let G be any graph.*

- (i) *The bonds of G , finite or infinite, are the circuits of a matroid $M_B(G)$, the bond matroid of G .*
- (ii) *The bond matroid of G is the dual of its finite-cycle matroid $M_{FC}(G)$.*

Proof. For simplicity we assume that G is connected; the general case is very similar. From [3] we know that $M_{FC}(G)$ has a dual; let us call this dual $M_B(G)$, and show that its circuits are the bonds of G . By definition of matroid duality, the circuits of $M_B(G)$ are the minimal edges sets that meet every spanning tree of G .

We show first that every bond B of G is a circuit of $M_B(G)$, a minimal set of edges meeting every spanning tree. Since B is a non-empty cut, it is the set of edges across some partition of the vertex set of G . Every spanning tree meets both sides of this partition, so it has an edge in B . On the other hand, we can extend any edge $e \in B$ to a spanning tree of G that contains no further from B , since by the minimality of B as a cut its two sides are connected in G . Hence B is minimal with the property of meeting every spanning tree.

Conversely, let B be any set of edges that is minimal with the property of meeting every spanning tree. We show that B contains a bond; by the implication already shown, and its minimality, it will then be that bond. Since G has a spanning tree, we have $B \neq \emptyset$; let $e \in B$. If B contains no bond, then every bond has an edge not in B . The subgraph H formed by all these edges is connected and spanning in G , as otherwise the edges of G from the component C of H containing e to any fixed component of $G - C$ would form a bond of G with no edge in H , contradicting its definition. So H contains a spanning tree. This misses B , contradicting the choice of B . \square

We prove Theorem 2 for countable graphs; the proof for arbitrary graphs can be deduced from this by considering a quotient space of $\|G\|$ as explained in [16]. For the remainder of this section, let G be a fixed countable connected graph.

We shall call two points in $\|G\|$ (*topologically*) *indistinguishable* if they have the same open neighbourhoods. Clearly two vertices or edge-ends $x, y \in \|G\|$ are indistinguishable if they cannot be separated by finitely many edges. (If both are edge-ends, then $x = y$.) On the other hand, two such points that can be separated by finitely many edges have disjoint open neighbourhoods. Inner points of edges are always distinguishable from all other points.

We shall need a few lemmas. Some of these are quoted from Schulz [16]; the others are adaptations of results proved in [10] for the special case that G is finitely separable. We remark that it is also possible to reduce Theorem 2 formally to that case by replacing G with a quotient graph as explained in [16].

Lemma 8. [16] *$\|G\|$ is a compact space.*

Lemma 9. *Let $X \subseteq \|G\|$ be a closed subspace, with disjoint open subsets O_1, O_2 such that $X = O_1 \cup O_2$. Then the set F of edges of G with one endvertex in O_1 and the other in O_2 is finite.*

Proof. Suppose F is infinite. As $\|G\|$ is compact, the mid-points of the edges in F have an accumulation point x in $\|G\|$. By definition of $\|G\|$, every neighbourhood of x contains infinitely many edges from F , and hence meets both O_1 and O_2 . Since O_1 and O_2 are closed in X , and hence in $\|G\|$, this means that $x \in O_1 \cap O_2$, contradicting our assumption that $O_1 \cap O_2 = \emptyset$. \square

In a Hausdorff space, every topological x - y path contains an injective such path, an x - y arc. Since $\|G\|$ is not necessarily Hausdorff we cannot assume this shortcut lemma in general, but it holds in the relevant case:

Lemma 10. [16] *If two points $x, y \in V(G) \cup \mathcal{E}(G)$ are separated by a finite set of edges, then every topological x - y path contains an x - y arc.*

Lemma 11. [16] *Let $x, y \in V(G) \cup \mathcal{E}(G)$, and let $(A_\gamma)_{\gamma < \lambda}$ be a transfinite sequence of x - y arcs in $\|G\|$. Then there exists a topological x - y path P and a dense subset P^* of P so that for all $p \in P^*$ the arcs A_γ containing p form a cofinal subsequence.*

Lemma 12. *Every closed connected subspace X of $\|G\|$ is path-connected.*

Proof. Suppose X is connected but not path-connected. Then there are $x, y \in V(G) \cup \mathcal{E}(G)$ contained in different path-components. In particular, x and y are topologically distinguishable, so they are separated by finitely many edges. Let e_1, e_2, \dots be a (possibly finite) enumeration of the edges in $E(G) \setminus E(X)$, let $F_i := \{e_1, \dots, e_i\}$ for all i . If there exists an i such that x and y lie in the closures of different graph-theoretical components of $G - F_i$, then picking an inner point outside X from every edge in F_i we obtain a finite set $Z \subseteq \|G\| \setminus X$ witnessing that x and y lie in distinct open sets of X whose union is all of X , contradicting our assumption that X is connected.

Hence for every i the points x and y lie in the closure $\overline{C_i}$ of the same component C_i of $G - F_i$. So for each i there is a path, ray or double ray connecting x to y in $\overline{C_i}$, and with Lemma 10 we then obtain an x - y arc A_i in $\overline{C_i}$. By Lemma 11 this implies that there is a topological x - y path P and a dense subset $P^* \subseteq P$ such that for every $p \in P^*$ the arcs A_i containing p form a cofinal subsequence. Suppose there exists a j such that $\dot{e}_j \subseteq P$. Then there must be a point $p \in \dot{e}_j \cap P^*$. However, none of the A_i with $i \geq j$ contains \dot{e}_j . Thus, P does not use any edge outside X . As X is closed, this implies that $P \subseteq X$. The required x - y arc in X can be found inside P by Lemma 10. \square

Lemma 13. *Let $F \subseteq E(G)$ be a set of edges whose closure in $\|G\|$ contains no circle. Then G has a topological spanning tree whose edge set contains F .*

Proof. Let $G = (V, E)$, let e_1, e_2, \dots be an enumeration of the edges in $E \setminus F$, and set $T_0 := E$. Inductively, if the closure of $(V, T_{i-1} - e_i)$ is connected in $\|G\|$ then set $T_i := T_{i-1} - e_i$; otherwise put $T_i := T_{i-1}$. Finally, we set $T := \bigcap_{i=0}^{\infty} T_i$.

In order to show that T is the edge set of a topological spanning tree, let us first check that the closure X of (V, T) is connected. Suppose there are two disjoint non-empty open sets O_1 and O_2 of X with $X = O_1 \cup O_2$. Then Lemma 9 implies that the cut S consisting of the edges with one endvertex in O_1 and the other in O_2 is finite. If j is the largest integer with $e_j \in S$ then, however, the closure of (V, T_j) is not connected, a contradiction. Thus, $\overline{T} = X$ is connected and therefore spanning. Moreover, \overline{T} is path-connected, by Lemma 12.

Secondly, we need to show that \overline{T} is acirclic. So, suppose that \overline{T} contains a circle C . Since every circle lies in the closure of its edges but the closure of $\bigcup F$ contains no circle, $E(C) \setminus F$ is non-empty. Pick j minimal with $e_j \in E(C) \setminus F$. Since e_j was not deleted from T_{j-1} when T_j was formed, the closure Y of $(V, T_{j-1} - e_j)$ is disconnected. So there are two disjoint non-empty open subsets O_1, O_2 of Y such that $Y = O_1 \cup O_2$. The endvertices of e_j do not lie in the same O_i , since adding e_j to that O_i would then yield a similar decomposition of the closure of (V, T_{j-1}) , contradicting its connectedness. But now the connected subset $C \setminus e_j$ of Y meets both O_1 and O_2 , a contradiction. Thus, \overline{T} does not contain any circle and is therefore a topological spanning tree. \square

Lemma 14. *Let C_1 and C_2 be two circles in $\|G\|$. Then $E(C_2) \subseteq E(C_1)$ implies that $E(C_1) = E(C_2)$.*

Proof. We first prove the following:

$$\text{For every point } x \in \overline{C_1} \setminus C_1 \text{ there is a point } y \in C_1 \text{ such that } x \text{ and } y \text{ are indistinguishable.} \quad (1)$$

Indeed, consider a $z \in \|G\|$ that is distinguishable from all points in C_1 . Thus, we may pick for every $p \in C_1$ two disjoint open neighbourhoods O_z^p and O_p of z and p , respectively. Note that C_1 is compact, being a continuous image of the compact space S^1 . Thus, there is a finite subcover $O_{p_1} \cup \dots \cup O_{p_n}$ of C_1 . Then, the open set $\bigcap_{i=1}^n O_z^{p_i}$ is disjoint from C_1 and contains z . Hence, z does not lie in the closure of C_1 . This proves (1).

Next, suppose that $E(C_2)$ is a proper subset of $E(C_1)$, and pick $e \in E(C_2)$ and $f \in E(C_1) \setminus E(C_2)$. Since $X := C_1 \setminus (\mathring{e} \cup \mathring{f})$ is disconnected there exist two disjoint non-empty open sets O'_1 and O'_2 of X with $X = O'_1 \cup O'_2$. For $j = 1, 2$, denote by I_j the set of points x in $\|G\|$ for which there is a $y \in O'_j$ such that x and y are indistinguishable. Then $O_1 := O'_1 \cup I_1$ and $O_2 := O'_2 \cup I_2$ are disjoint and open subsets of $X \cup I_1 \cup I_2$. Moreover, it follows from (1) that $\overline{C_1} \setminus (\mathring{e} \cup \mathring{f}) = X \cup I_1 \cup I_2$. Therefore, O_1 and O_2 are two disjoint non-empty open sets of $\overline{C_1} \setminus (\mathring{e} \cup \mathring{f})$ with $\overline{C_1} \setminus (\mathring{e} \cup \mathring{f}) = O_1 \cup O_2$.

As $C_2 \setminus \mathring{e}$ is a connected subset of $\overline{C_1} \setminus (\mathring{e} \cup \mathring{f})$ it lies in O_1 or in O_2 , let us say in O_1 . Then $\tilde{O}_1 := O_1 \cup \mathring{e}$ and O_2 are two disjoint non-empty open subsets of $\overline{C_1} \setminus \mathring{f}$ with $\overline{C_1} \setminus \mathring{f} = \tilde{O}_1 \cup O_2$. By (1), this means that also $C_1 \setminus \mathring{f}$ is disconnected. But $C_1 \setminus \mathring{f}$ is a continuous image of a connected space, and hence connected. \square

Lemma 15. *Let T be a standard subspace of $\|G\|$. Then the following statements are equivalent:*

- (i) *T is a topological spanning tree of $\|G\|$.*
- (ii) *T is maximally acirclic, that is, it does not contain a circle but adding any edge in $E(G) \setminus E(T)$ creates one.*
- (iii) *$E(T)$ meets every finite bond, and is minimal with this property.*

Proof. Let us first prove a part of (iii) \rightarrow (i) before dealing with all the other implications.

If $E(T)$ meets every finite bond then T is spanning and path-connected. (2)

Suppose that the closure X of $(V(G), E(T))$ is not connected. Then there are two disjoint non-empty open sets O_1 and O_2 of X with $X = O_1 \cup O_2$. From Lemma 9 we get that the cut consisting of the edges with one endvertex in O_1 and the other in O_2 is finite. Since each of O_1 and O_2 needs to contain a vertex, this cut is non-empty. Hence, $E(T)$ misses a finite bond, a contradiction. Therefore, $T = X$ is connected and then, by Lemma 12, path-connected.

(i) \rightarrow (ii) Consider any edge $e \notin E(G) \setminus E(T)$. If the endvertices u and v of e cannot be separated by finitely many edges then $e - u$ (and also $e - v$) is a circle in $\|G\|$. Otherwise, any topological $u-v$ path contains an $u-v$ arc by Lemma 10. In particular, T contains an $u-v$ arc that together with e forms a circle.

(ii) \rightarrow (iii) Suppose that $E(T)$ misses a finite bond F . Pick $e \in F$, and let C be a circle in $T \cup e$ through \dot{e} . Pick an inner point of every edge in F and denote the set of these points by Z . Then the two components of $\|G\| \setminus Z$, each of which contains an endvertex of e , form two disjoint open sets containing T . However, $C \setminus \dot{e} \subseteq T$ is a connected set that meets both of these disjoint open sets, which is impossible. Thus, $E(T)$ meets every finite bond. In particular, T is spanning and path-connected, by (2).

Let f be any edge in $E(T)$, and let us show that $E(T) - f$ misses some finite bond. Denote the endvertices of f by r and s , and observe that r and s can be separated by finitely many edges as T is acirclic. Denote by K_r and K_s the path-components of $T \setminus \dot{f}$ containing r and s , respectively. By Lemma 10 and as T does not contain any circle, K_r and K_s are distinct, and thus disjoint. As T is path-connected, it follows that $T \setminus \dot{f}$ is the disjoint union of the open sets K_r and K_s . Now Lemma 9 yields that there are only finitely many edges with one endvertex in K_r and the other in K_s . As T is spanning this means that $E(T) - f$ misses a finite cut.

(iii) \rightarrow (i) By (2), we only need to check that T does not contain any circle. Suppose there exists a circle $C \subseteq T$, and pick some $e \in E(C)$. By the minimality of $E(T)$ there exists a finite bond F so that F is disjoint from $T \setminus \dot{e}$. Then, however, picking inner points from the edges in F yields a set Z , so that the connected set $C \setminus \dot{e}$ is contained in $\|G\| \setminus Z$ but meets two components of $\|G\| \setminus Z$, which is impossible. \square

We can now prove our main theorem, which we restate:

Theorem 2.

- (i) *The edge sets of the circles in $\|G\|$ are the circuits of a matroid $M_C(G)$, the cycle matroid of G .*
- (ii) *The bases of $M_C(G)$ are the edge sets of the topological spanning trees of G .*
- (iii) *The cycle matroid $M_C(G)$ is the dual of the finite-bond matroid $M_{FB}(G)$.*

Proof. To bypass the need to verify any matroid axioms, we *define* $M_C(G)$ as the dual of $M_{FB}(G)$ (which we know exists [3]), ie., as the matroid whose bases B are the complements of the bases of $M_{FB}(G)$. These latter are the maximal edge sets not containing a finite bond, so the bases B of $M_C(G)$ are the minimal edge sets meeting every finite bond. By Lemma 15 below, this is equivalent to B being the edge set of a topological spanning tree of $\|G\|$.

We have defined $M_C(G)$ so as to make (iii) true, and shown (ii). It remains to show (i): that the circuits of $M_C(G)$ are the edge sets of the circles in $\|G\|$. Since no circuit of a matroid contains another circuit, and since by Lemma 14 no edge set of a circle contains another such set, it suffices to show that every circuit contains the edge set of a circle, and conversely every edge set of a circle contains a circuit.

For the first of these statements note that, by assertion (ii), a circuit D of $M_C(G)$ does not extend to the edge set of a topological spanning tree. Hence by Lemma 13 its closure $\bigcup D$ in $\|G\|$ contains a circle C . For the second statement, note that the edge set D of a circle C is not contained in the edge set of a topological spanning tree T , because T is closed and would therefore contain $\bigcup D \supseteq C$, contradicting its definition. Hence D is dependent in $M_C(G)$, and therefore contains a circuit [3]. \square

Finally, let us restate and prove Theorem 5:

Theorem 5. *The following three assertions are equivalent for a countable finitely separable graph G :*

- (i) *G is planar;*
- (ii) *$M_C^*(G)$ is finitely graphic;*
- (iii) *$M_{FC}^*(G)$ is graphic.*

Proof. If G is planar, it has a dual G^* . Then $M_C^*(G) = M_{FC}(G^*)$ by Theorem 4, and $M_{FC}^*(G) = M_B(G) = M_C(G^*)$ by Theorems 1 and 3 (ii).

(ii) \rightarrow (i): Since $M_C^*(G)$ is finitely graphic, there exists a graph H with the same edge set as G such that $M_C^*(G) = M_{FC}(H)$. As $M_{FC}^*(H) = M_B(H)$ by Theorem 1 and matroid duals are unique, we obtain $M_C(G) = M_B(H)$. Hence the edge sets of the circles in $\|G\|$, which by Theorem 2 (i) are the circuits

of $M_C(G)$, are precisely the bonds of H . So H is a dual of G , and G is planar by Theorem 3 (i).

(iii)→(i): Let H be a graph such that $M_C(H) = M_{FC}^*(G)$. To show that G is planar, it suffices by Kuratowski's theorem¹⁰ to check that G has no K_5 - or $K_{3,3}$ -minor, or in matroid terms, that $M_{FC}(G)$ has no minor isomorphic to $M_C(K_5)$ or $M_C(K_{3,3})$. As $M_{FC}^*(G) = M_C(H)$, this is equivalent to saying that $M_C(H)$ has no $M_C^*(K_5)$ or $M_C^*(K_{3,3})$ -minor. These latter two matroids are not graphic [14, Prop. 2.3.3], so it thus suffices to show that the finite minors of $M_C(H)$ are graphic.

To prove this, consider a finite minor M of $M_C(H)$, obtained by deleting the set $X \subseteq E(H)$ and contracting the set $Y \subseteq E(H)$, say. Let V be the finite set of vertices of H incident with an edge in the ground set E of M , and let K be the finite graph obtained from the graph (V, E) by identifying any two vertices that are either indistinguishable in $\|H\|$ or joined by an arc in $\|H\|$ whose edges lie in some fixed base B of the restriction of $M_C(H)$ to Y . Using Lemma 10 and [3, Lemma 3.5], it is now easy to show that $M = M_C(K)$. Hence M is graphic, as desired. \square

We remark that the graphs witnessing (ii) and (iii) in Theorem 5 can be chosen to be finitely separable, too. Indeed, the graph G^* which we used in our proof as a witness for both (ii) and (iii) is finitely separable by Theorem 3 (ii).

5 The dual of the algebraic cycle matroid

Recall from Theorem 6 that the elementary algebraic cycles of a graph G are the circuits of a matroid, the *algebraic cycle matroid* $M_{AC}(G)$ of G , if and only if G contains no subdivision of the graph shown in Figure 1. In this section we characterize their matroid $M_{AC}^*(G)$.

Recall that a *ray* is a 1-way infinite path. Let us call a non-empty cut $F = E(A, B)$ of G *skew* if one of its sides A, B is *small* in the sense that the subgraph it induces in G contains no ray and F is minimal with this property among the non-empty cuts of G . If G is connected then so is the small side of any skew cut, so this will be finite if G is connected and locally finite.

Theorem 16. *The cocircuits of a matroid M that is the algebraic cycle matroid of a graph G are precisely the skew cuts of G .*

Casteels and Richter [6] studied a related duality: they showed that, in a locally finite graph, the cuts with a finite side form the orthogonal space of the set of elementary algebraic cycles.

For our proof of Theorem 16 we need the following easy lemma from [3]:

Lemma 17. *A circuit and a cocircuit of a matroid never meet in exactly one element.*

¹⁰Its extension to countable graphs is straightforward by compactness; see [9, Exercise 8.23].

Proof of Theorem 16. Let us show first that every non-empty cut F with a small side A contains a cocircuit of M . If not, F is independent in M^* , so it avoids a base B of M . Adding an element $f \in F$ to B creates a circuit of M , a cycle or double ray that meets F precisely in f . Since $G[A]$ contains no ray, this is impossible.

Conversely, let us show that every cocircuit D of M contains a skew cut. Pick an edge $e \in D$. If its endvertices lie in the same component of $G - D$, then G contains a cycle meeting D in exactly e , contradicting Lemma 17. So the endvertices of e lie in distinct components of $G - D$. If both these contain a ray, then these rays can be chosen so as to combine with e to a double ray meeting D precisely in e , again contradicting Lemma 17. Hence one of these components contains no ray. Its vertex set A is the small side of a cut F with $e \in F \subseteq D$. To show that F is a skew cut, we still have to show that it contains no non-empty cut F' with a small side properly. But any such F' contains a cocircuit D' , as shown earlier, giving $D' \subseteq F' \subsetneq F \subseteq D$. This contradicts the fact that no cocircuit contains another cocircuit properly.

We have shown that every skew cut contains a cocircuit, and vice versa. Since skew cuts, as cocircuits, cannot contain each other properly, these inclusions cannot be proper. So the cocircuits of M are the skew cuts of G . \square

6 Thin-sum matroids and representability

In finite matroid theory, representable matroids are an important generalization of graphic matroids. As matroids defined by linear independence are finitary, the dual of an infinite representable matroid will not, except in trivial cases, be representable. Representability, as usually defined, is thus another concept that seems too narrow for infinite matroids.

In this section we present a notion of vector independence, different from linear independence, that can give rise to non-finitary matroids. Examples include the algebraic cycle matroids of graphs and of higher-dimensional complexes [3].

Let F be a field, and let A be some set. We say that a set X of functions $x : A \rightarrow F$ is *thin* if for every $a \in A$ there are only finitely many $x \in X$ with $x(a) \neq 0$. Given such a thin set of functions, their pointwise sum $\sum_{x \in X} x$ is another $A \rightarrow F$ function. We say that a set of $A \rightarrow F$ functions, not necessarily thin,¹¹ is *thinly independent* if for every thin subset X and every corresponding family $(\alpha_x)_{x \in X}$ of coefficients $\alpha_x \in F$ we have $\sum_{x \in X} \alpha_x x = \mathbf{0} \in F^A$ only when $\alpha_x = 0$ for all $x \in X$.

Unlike linear independence, thin independence does not always define a matroid.¹² The following theorem gives a sufficient condition for when it does:

¹¹Requiring independent sets to be thin leads to a different notion of representability that may have its own applications. It is easily seen that this notion does not satisfy (IM) for arbitrary sets E of $A \rightarrow F$ functions, but there may be interesting examples where it does.

¹²View the elements of $E = \mathbb{F}_2^{\mathbb{N}}$ as subsets of \mathbb{N} , and define sets $I := \{\{1, n\} : n \in \mathbb{N}\}$ and $I' := \{\{n\} : n \in \mathbb{N}\}$. Both I and I' are thinly independent. Moreover, I' is maximally thinly independent but I is not: $I + \mathbb{N}$, for instance, is still thinly independent. Yet, the only $x \in I'$ for which $I + x$ is thinly independent is $x = \{1\}$, which is already contained in I . Thus, (I3)

Theorem 18. *If a set E of $A \rightarrow F$ functions is thin, then its thinly independent subsets form the independent sets of a matroid on E .*

Let us call such a matroid as in Theorem 18 the *thin-sums matroid* of the functions in E . In the remainder of this section we prove Theorem 18, and then briefly discuss what it means for graphs.

Let F be endowed with the discrete topology, and the set F^A of all $A \rightarrow F$ functions with the product topology. Thus for each $x \in F^A$, the sets

$$\{y \in F^A : y(a) = x(a) \text{ for all } a \in A'\}$$

where A' ranges over the finite subsets of A forms a basis of the open neighbourhoods of x .

Given a set X of functions $A \rightarrow F$, we write $\langle X \rangle$ for the set of all functions $A \rightarrow F$ that are of the form $\sum_{x \in X'} \alpha_x x$, where $X' \subseteq X$ and X' is thin. Similarly, we write \overline{X} for the closure in F^A of the set X .

In contrast to [4] where these concepts were introduced, X will here always be a subset of a thin set E . While the sets $\langle X \rangle$ and \overline{X} may contain elements outside our ground set E , we note that $X \mapsto \langle X \rangle \cap E$ is the closure operator associated with the set \mathcal{I} of thinly independent subsets of E , as defined in Section 2.

We need two lemmas from [4], which together imply that $\langle \langle X \rangle \rangle = \langle X \rangle$ for all our (thin) sets $X \subseteq E$:

Lemma 19. [4, Lemma 5] *Every thin set $X \subseteq F^A$ satisfies $\overline{\langle X \rangle} = \langle X \rangle$.*

Lemma 20. [4, Lemma 6] *Every set $X \subseteq F^A$ satisfies $\langle \langle X \rangle \rangle \subseteq \overline{\langle X \rangle}$.*

Proof of Theorem 18. Let \mathcal{I} be the set of thinly independent subsets of E . Clearly, \mathcal{I} satisfies (I1) and (I2).¹³

Our first claim is that, for all sets $J \subseteq X \subseteq E$ with $J \in \mathcal{I}$,

$$\text{if } X \subseteq \langle J \rangle \text{ then } J \text{ is a maximal element of } \{I \in \mathcal{I} : I \subseteq X\}. \quad (3)$$

Consider an $x \in X \setminus J$. Then there are coefficients $a_j \in F$, $j \in J$, such that $\sum_{j \in J} a_j j = x$. Thus, $J + x$ is not thinly independent for any $x \in X \setminus J$, which implies the claim.

To prove that \mathcal{I} satisfies (I3), let $I \in \mathcal{I} \setminus \mathcal{I}^{\max}$ and $I' \in \mathcal{I}^{\max}$ be given. We have to find an $x \in I' \setminus I$ such that $I + x$ is still thinly independent. Clearly, any x in $I' \setminus \langle I \rangle$ will do, so it suffices to show that $I' \not\subseteq \langle I \rangle$. If $I' \subseteq \langle I \rangle$, then $\langle I' \rangle \subseteq \langle \langle I \rangle \rangle = \langle I \rangle$ by Lemmas 19 and 20. As $I' \in \mathcal{I}^{\max}$ implies $E \subseteq \langle I' \rangle$, this yields $E \subseteq \langle I \rangle$, which contradicts (3). This completes the proof of (I3).

is violated.

¹³We use the independence axioms in our proof. Alternatively, one could check that $\langle \cdot \rangle \cap E$ is indeed the closure operator associated with \mathcal{I} and then use the closure axioms: (CL1), (CL2) and (CL4) are straightforward, (CL3) follows from our two lemmas, and (CLM) is proved like (IM) in the text.

Next, we prove that \mathcal{I} satisfies (IM). This will follow directly from (3) and the following claim:

For all sets $I \subseteq X \subseteq E$ with $I \in \mathcal{I}$, there is a $B \in \mathcal{I}$ with $I \subseteq B \subseteq X$ and $X \subseteq \langle B \rangle$. (4)

In the remainder of this proof we thus prove (4). Let x_1, x_2, \dots be a (possibly transfinite) enumeration of $X \setminus I$. Inductively, we define nested sets $B_\lambda \subseteq X$ as follows. Start with $B_0 := I$. If, in step λ , there are families $(\alpha_\mu)_{\mu > \lambda}$ and $(\beta_i)_{i \in I}$ of coefficients in F such that

$$x_\lambda = \sum_{\mu > \lambda} \alpha_\mu x_\mu + \sum_{i \in I} \beta_i i$$

(these sums are well-defined, since E is thin by assumption), we set $B_\lambda := \bigcup_{\mu < \lambda} B_\mu$. Otherwise we put $B_\lambda := \{x_\lambda\} \cup \bigcup_{\mu < \lambda} B_\mu$. Finally, we let $B := \bigcup_\lambda B_\lambda$, which we claim satisfies (4).

Let us first check that $B \in \mathcal{I}$. If not then there are coefficients $\alpha_b \in F$ for all $b \in B$, not all of them zero, such that $\sum_{b \in B} \alpha_b b = \mathbf{0}$. Since I is thinly independent, there must be some $b \in B \setminus I$ with $\alpha_b \neq 0$. Pick such a function $b = x_\lambda$ with smallest index λ . Then

$$-x_\lambda = \sum_{x \in \{x_\mu : \mu > \lambda\} \cap B} \alpha_{x_\lambda}^{-1} \alpha_x x + \sum_{i \in I} \alpha_{x_\lambda}^{-1} \alpha_i i,$$

which contradicts the fact that $x_\lambda \in B$.

Next, we prove $X \subseteq \overline{\langle B \rangle}$. By Lemma 19 this will imply $X \subseteq \overline{\langle B \rangle} = \langle B \rangle$, which then completes the proof of (4).

To prove $X \subseteq \overline{\langle B \rangle}$, consider a function $z \in X$. We need to find, for every finite subset A' of A , a function $z' \in \langle B \rangle$ that agrees with z on A' . Denote by L the set of $x \in X$ for which there is an $a \in A'$ with $x(a) \neq 0$. Observe that L is a finite set, since A' is finite and $X \subseteq E$ is thin. In particular, we may write $L \setminus B = \{x_{\lambda_1}, \dots, x_{\lambda_k}\}$, where $\lambda_1 < \dots < \lambda_k$. Let ℓ be the largest number in $\{1, \dots, k+1\}$ for which there exists a $y \in \langle \{x_{\lambda_\ell}, \dots, x_{\lambda_k}\} \cup (B \cap L) \rangle$ with $y(a) = z(a)$ for all $a \in A'$. Note that there is always such a y for $\ell = 1$, as we may either pick $y = z$ if $z \in L$, or $y = \mathbf{0}$ otherwise. Note, furthermore, that we have found the desired z' if $\ell = k+1$, as then $z' := y \in \langle B \rangle$.

Suppose that $\ell \leq k$. Since $x_{\lambda_\ell} \notin B$ there are coefficients $(\alpha_\mu)_{\mu > \lambda_\ell}$ and $(\beta_i)_{i \in I}$ such that $x_{\lambda_\ell} = \sum_{\mu > \lambda_\ell} \alpha_\mu x_\mu + \sum_{i \in I} \beta_i i$. Then

$$x_{\lambda_\ell} = \sum_{\mu > \lambda_\ell, x_\mu \notin B} \alpha_\mu x_\mu + \sum_{b \in B} \beta'_b b$$

with suitable coefficients β'_b . Restricting this to L , set

$$r := \sum_{p=\ell+1}^k \alpha_{\lambda_p} x_{\lambda_p} + \sum_{b \in B \cap L} \beta'_b b.$$

Then

$$r(a) = x_{\lambda_\ell}(a) \text{ for all } a \in A', \text{ and } r \in \langle \{x_{\lambda_{\ell+1}}, \dots, x_{\lambda_k}\} \cup (B \cap L) \rangle.$$

Next, by choice of ℓ there exists $y = \sum_{q=\ell}^k \gamma_{\lambda_q} x_{\lambda_q} + \sum_{b \in B \cap L} \delta_b b$ such that $y(a) = z(a)$ for all $a \in A'$. Replacing x_{λ_ℓ} in this sum with r , we obtain

$$y^* := \gamma_{\lambda_\ell} r + \sum_{q=\ell+1}^k \gamma_{\lambda_q} x_{\lambda_q} + \sum_{b \in B \cap L} \delta_b b \in \langle \{x_{\lambda_{\ell+1}}, \dots, x_{\lambda_k}\} \cup (B \cap L) \rangle.$$

As y^* agrees with z on A' , this contradicts the maximal choice of ℓ . \square

The algebraic cycle matroid of a graph $G = (V, E)$ can be represented as a thin-sums matroid over \mathbb{F}_2 , for any G for which it is defined (cf. Theorem 6). Indeed, as in finite graphs we represent an edge $e = uv$ by the map $V \rightarrow \mathbb{F}_2$ assigning 1 to both u and v , and 0 to every other vertex. Then a set $F \subseteq E$ of edges becomes a set of $V \rightarrow \mathbb{F}_2$ functions, not necessarily thin, which is thinly independent if and only if F contains no elementary algebraic cycle. This example can be generalized to higher dimensions; see [3] for algebraic cycle matroids of simplicial complexes.

We do not know whether the other non-finitary matroids we discussed in this paper are representable as thin-sum matroids, but suspect not. For finitary matroids, one would hope that ‘thin-sum’ representability coincides with traditional representability, but we have no proof of this:

Problem 21. *Is a finitary matroid representable as a thin-sums matroid if and only if it is representable in the usual sense?*

We can show that the finite-cycle matroid $M_{\text{FC}}(G)$ of a graph G is a thin-sums matroid if G has finite chromatic number, but we do not know this for arbitrary G .

Acknowledgement

We would like to thank Robin Christian for indicating the proof of implication (iii) \rightarrow (i) of Theorem 5, which we had previously stated as a conjecture.

References

- [1] D.W.T. Bean, *A connected finitary co-finitary matroid is finite*, Proceedings of the Seventh Southeastern Conference on Combinatorics, Graph Theory and Computing, Congressus Numerantium, vol. 17, 1976, pp. 115–19.
- [2] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. Comput. **15** (2006), 75–90.

- [3] H. Bruhn, R. Diestel, M. Kriesell, R. Pendavingh, and P. Wollan, *Axioms for infinite matroids*, arXiv:1003.3919 (2010).
- [4] H. Bruhn and A. Georgakopoulos, *Bases and closures under infinite sums*, Linear Algebra and its Applications **435** (2011), 2007–2018.
- [5] H. Bruhn and P. Wollan, *Finite connectivity in infinite matroids*, J. Combin. Theory (Series B) (to appear).
- [6] K. Casteels and B. Richter, *The Bond and Cycle Spaces of an Infinite Graph*, J. Graph Theory **59** (2008), no. 2, 162–176.
- [7] R. Christian, R.B. Richter, and B. Rooney, *The planarity theorems of MacLane and Whitney for graph-like continua*, Electronic J. Comb. **17** (2009), #R12.
- [8] R. Diestel, *Locally finite graphs with ends: a topological approach*, <http://arxiv.org/abs/0912.4213>, 2009.
- [9] ———, *Graph Theory*, 4th ed., Springer, 2010.
- [10] R. Diestel and D. Kühn, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.
- [11] D.A. Higgs, *Infinite graphs and matroids*, Recent Progress in Combinatorics, Proceedings Third Waterloo Conference on Combinatorics, Academic Press, 1969, pp. 245–53.
- [12] ———, *Matroids and duality*, Colloq. Math. **20** (1969), 215–220.
- [13] J.G. Oxley, *Infinite matroids*, Matroid applications (N. White, ed.), Encycl. Math. Appl., vol. 40, Cambridge University Press, 1992, pp. 73–90.
- [14] ———, *Matroid theory*, Oxford University Press, 1992.
- [15] R. Rado, *Abstract linear dependence*, Colloq. Math. **14** (1966), 257–64.
- [16] M. Schulz, *Der Zyklenraum nicht lokal-endlicher Graphen*, Diplomarbeit Universität Hamburg 2005, <http://www.math.uni-hamburg.de/home/diestel/papers/others/Schulz.Diplomarbeit.pdf>.
- [17] M. Stein, *Extremal infinite graph theory*, Discrete Math. **311** (2011), 1472–1496.
- [18] C. Thomassen, *Duality of infinite graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160.
- [19] M. Las Vergnas, *Sur la dualité en théorie des matroïdes*, Théorie des Matroïdes, Lecture notes in mathematics, vol. 211, Springer-Verlag, 1971, pp. 67–85.

Version 12 March 2011.

Footnotes 5 and 6 were added later. Figure 1 was corrected in May 2012.